ARMY RESEARCH LABORATORY

Exploiting Social Media for Army Operations: Syrian Civil War Use Case

by Sue E. Kase, Elizabeth K. Bowman, Tanvir Al Amin, and Tarek Abdelzaher

ARL-RP-0489 **July 2014**

A reprint from *Proc. of SPIE Vol. 9122, 91220D, 2014*

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory
Aberdeen Proving Ground, MD 21005

ARL-RP-0489 **July 2014**

Exploiting Social Media for Army Operations: Syrian Civil War Use Case

Sue E. Kase, Elizabeth K. Bowman, Tanvir Al Amin, and Tarek Abdelzaher
Computational and Information Sciences Directorate, ARL

A reprint from *Proc. of SPIE Vol. 9122, 91220D, 2014*

REPORT DOCUMENTATION PAGE	Form Approved OMB No. 0704-0188

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
July 2014	Reprint	

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Exploiting Social Media for Army Operations: Syrian Civil War Use Case	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)	5d. PROJECT NUMBER
Sue E. Kase, Elizabeth K. Bowman, Tanvir Al Amin, and Tarek Abdelzaher	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
U.S. Army Research Laboratory ATTN: RDRL-CII-C Aberdeen Proving Ground, MD 20783-1197	ARL-RP-0489

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Millions of people exchange user-generated information through online social media (SM) services. The prevalence of SM use globally and its growing significance to the evolution of events has attracted the attention of many agencies, from humanitarian non-government organizations (NGOs) and disaster response agencies to homeland security and counter-terrorism. The information exchanged in SM sites and the networks of people who interact with these online communities can provide insights into ongoing events. For example, SM could provide ongoing assessment of disaster relief and humanitarian operations from a local perspective, or offer an understanding of risk levels to which the operators in question are exposed. Despite its potential value, there are significant technological barriers to leveraging SM. SM collection and analysis are difficult in the dynamic SM environment and deception is a real concern. This paper introduces a credibility analysis approach and prototype fact-finding technology called the "Apollo Fact-finder" that mitigates the problem of inaccurate or falsified SM data. Apollo groups data into sets (or claims), corroborating specific observations, then iteratively assesses both claim and source credibility resulting in a ranking of claims by likelihood of occurrence. These credibility analysis approaches are discussed in the context of a hypothetical humanitarian mission executed in an area of active conflict and applied to public domain tweets collected in the aftermath of a Syrian crisis.

15. SUBJECT TERMS

Social media, Text analytics, Credibility analysis

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
A. Report	b. ABSTRACT	c. THIS PAGE			Sue E. Kase
Unclassified	Unclassified	Unclassified	UU	18	19b. TELEPHONE NUMBER *Include area code* (410) 278-9762

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Exploiting social media for Army operations: Syrian crisis use case

Sue E. Kase*[a], Elizabeth K. Bowman[a], Md Tanvir Al Amin[b], Tarek Abdelzaher[b]
[a]Army Research Laboratory, Computational and Information Science Directorate, Tactical Information Fusion Branch, Aberdeen Proving Ground, MD USA;
[b]Cyber Physical Computing Laboratory, Department of Computer Science, University of Illinois Urbana Champaign, Urbana, IL USA

ABSTRACT

Millions of people exchange user-generated information through online social media (SM) services. The prevalence of SM use globally and its growing significance to the evolution of events has attracted the attention of many agencies, from humanitarian non-government organizations (NGOs) and disaster response agencies to homeland security and counter-terrorism. The information exchanged in SM sites and the networks of people who interact with these online communities can provide insights into ongoing events. For example, SM could provide ongoing assessment of disaster relief and humanitarian operations from a local perspective, or offer an understanding of risk levels to which the operators in question are exposed. Despite its potential value, there are significant technological barriers to leveraging SM. SM collection and analysis are difficult in the dynamic SM environment and deception is a real concern. This paper introduces a credibility analysis approach and prototype fact-finding technology called the "Apollo Fact-finder" that mitigates the problem of inaccurate or falsified SM data. Apollo groups data into sets (or claims), corroborating specific observations, then iteratively assesses both claim and source credibility resulting in a ranking of claims by likelihood of occurrence. These credibility analysis approaches are discussed in the context of a hypothetical humanitarian mission executed in an area of active conflict and applied to public domain tweets collected in the aftermath of a Syrian crisis.

Keywords: Social media, Text analytics, Credibility analysis

1. INTRODUCTION

A distinctive feature of broadcast social networks such as Twitter is that each tweet is regarded as a global broadcast destined to all users (unless specifically restricted otherwise), making tremendous amounts of public information available to all. With numbers as staggering as 400 million public tweets per day[1] and over 700 million daily active users on Facebook[2], there is no argument that social media (SM) has emerged as the dominant form of communication in today's society. It has led to exponential growth in rates of information diffusion, or how rapidly and widely a new idea or action spreads through communication channels[3]. In fact, it has been shown that 'retweeted' tweets are read by an average of 1,000 users no matter what the number of followers of the original tweet or user is[4].

With these statistics, it is no wonder that SM played a role in accelerating the Middle East uprisings that started in the first quarter of 2011. According to a Cairo activist: "*We used Facebook to schedule the protests, Twitter to coordinate, and YouTube to tell the world*"[5]. The power of SM is put into perspective when one considers that after a mere 18 days of continuous SM-driven protests from January 25, 2011 to February 11, 2011 Egypt's President Hosni Mubarak's thirty year regime came to an abrupt end. This illustrates change can be lighting fast, with SM serving as the engine behind that change, and with the results having profound consequences.

In a host-nation's densely populated areas comprised of locals with varying hostile, neutral, and friendly attitudes and allegiances, the local populace's use of SM to organize, schedule, and communicate represents a critical stream of information that is publically broadcast and could help optimize execution of missions, such as, for example, peace keeping operations and delivery of humanitarian aid. In Section 2 we discuss SM in the context of the Army and the role it might play in future theatres. We note that SM offers opportunities to study and better understand host-nation human terrain, but not without significant risks. Section 3 situates SM in the context of a conflict zone where pro- and anti-regimes created partial, misleading, and motivated narratives that need to be properly interpreted for risk assessment and mission support. Section 4 introduces a credibility analysis approach and prototype fact-finding technology applied to tweets collected in the aftermath of the Syrian humanitarian crisis. Section 5 discusses these analysis approaches from an Army perspective and suggests open challenges in the quest to automate the human credibility judgment process.

*sue.e.kase.civ@mail.mil; phone 1 410 278-9762

Next-Generation Analyst II, edited by Barbara D. Broome, David L. Hall, James Llinas,
Proc. of SPIE Vol. 9122, 91220D · © 2014 SPIE · CCC code: 0277-786X/14/$18
doi: 10.1117/12.2049701

2. SOCIAL MEDIA AND THE ARMY

An interesting open question to the research community is whether and how one could leverage and employ SM to support decision making and mission operations. SM's ostensible value to Army activities and operations stems largely from what it *might* provide and how it *might* be used[6].

This paper posits three propositions with regard to SM and Army needs. First, SM represents a growing and potentially valuable source of public information insofar as it houses and facilitates voluntary broadcast of information relevant to the day-to-day attitudes, behaviors, and intentions of social groups interested in activities that may impact government interests. Second, SM has the potential to facilitate a number of engagement and shaping endeavors as well as inform and influence direct action. Third, SM is unique in that it can be used to simultaneously disseminate information and evaluate the effectiveness of efforts through the same means and mechanisms. Clearly, the Army needs to understand SM and the opportunities it offers, as well as, the risks associated with its use.

2.1 Social Media Opportunities

SM can be a powerful tool for command-level decision makers, helping them understand and shape their Areas of Responsibility (AOR). The skillful leveraging and employment of SM can allow decision makers to favorably utilize communities, improve the quality and timeliness of the delivery of relevant information, and increase unity of effort in operations.

If decision makers make the effort to maintain, leverage, and employ a SM presence, they are better positioned to perceive emerging threats, but also to better understand the local human terrain[7]. SM is particularly valuable for places that are difficult to directly observe as it can let an analyst look in places that are not normally visible through other means[8]. SM can provide new data and information sources for increased situational awareness (SA) and understanding "the big picture." At the tactical and operational levels, this increased SA can help position commanders to make more informed and nuanced decisions about deploying assets or spending resources[9].

At higher levels, SM affords a chance to study and understand local culture and behavior that might otherwise be difficult to interpret. Local communities are a potential source of observers that far exceed the capabilities of deployed units, but more importantly, may possess linguistic, cultural, and local contextual knowledge and expertise upon which decision-makers can draw[10]. In this way, SM can help cover for capability gaps in language translation and intercultural communication. Local expertise goes beyond language and cultural gaps though, in that it provides a naturalistic way to distinguish what data and information are meaningful in the eyes of the community. This local perspective is important as it allows decision makers to not only better understand what is relevant and meaningful in a specific context, but also the low order effects of particular actions. This can make Course of Action generation, evaluation, and selection much more realistic and less uncertain.

2.2 Social Media Risks

As powerful as SM can be for understanding and shaping the AOR, there are important limitations on leveraging and employing SM. The ability to process "big data", that is also likely to be only semi-structured at best, and, depending on the AOR, in languages other than English, presents many technical barriers.

For all the promise of SM, the fact remains that SM-specific technology is still nascent. Existing off-the-shelf SM platforms have not been designed with Army operational considerations in mind, which means if the Army wants to leverage or employ SM for learning and interaction purposes, it has to make significant investments in purposefully designed platforms and software. Emerging technologies need to limit the volume of data through processing and analysis, as well as, verify the relevance and accuracy of the information[10].

The veracity of SM data and information is a particular problem for command-level decision makers because of the high stakes of taking action upon erroneous information. SM users can hide their identity and location leaving commanders unwilling to act on un-sourced, possibly inaccurate information[10]. Online identities may be deliberate fabrications ("sockpuppets") and SM may be deliberately seeded with false or biased information to aid opposed interests ("astroturfing")[11]. Sophisticated SM users are able to manipulate SM to distort perceptions or misrepresent public sentiment and discourse. In addition, SM information can be inaccurate because the user base is not representative of the larger discussion, or only representative of a particular locality with localized concerns[12].

There is no better example of the challenges associated with evaluating the credibility and significance of SM discourse than Syria's civil war—the most socially mediated civil conflict in history[13], and a target of many humanitarian efforts to remedy the side effects. An exceptional amount of what the outside world knows, or thinks it knows, about Syria's nearly three year old conflict has come from videos, analysis, and commentary circulated through SM.

3. SOCIAL MEDIA AND SYRIA'S CIVIL WAR

The uprisings in Tunisia and Egypt that transformed the Arab world inspired Syrian activists, who drew on similar methods used by other Arab activists across the region. Syrian activists posted videos to YouTube, adopted similar slogans ("the people want to overthrow the regime"), created Twitter hashtags (#mar15), and attempted to portray an image of a rising nonviolent Syrian protest wave through SM[13]. This impression did not necessarily reflect the reality on the ground at the time.

As the months progressed the balance within the opposition shifted toward armed groups. By the spring of 2012, Syria's conflict looked more like a civil war than the earlier peaceful uprising. The August 2012 resignation of UN special envoy Kofi Annan triggered a rapid cascade toward violent uprisings increasing the pace of devastation and displacement. As the protests became more dangerous, a pattern of media reliance on activist-generated online content was established in the absence of journalists present on the ground. The Syrian opposition crafted narratives for the international media of a peaceful, pro-Western uprising, while the Syrian regime sought to portray their opposition as radical Islamists supporting colluding outsiders.

SM proved essential to the international coverage of Syria from the outset. The nature of the Syrian regime and of the conflict meant limited direct access to the battlefields. Most television stations relied heavily on citizen journalists and online content for footage to accompany their stories. The sheer volume of information, videos, and discourse flowing from Syria in principle allowed the outside world unmediated access to the conflict in all its diversity. However, the politically motivated curation of SM combined with the deluge of information made it difficult to keep up or to evaluate the credibility and significance of information circulating online. As the battle between anti- and pro-regime groups accelerated, authentication and verification of information became increasingly important. It is against this backdrop that this paper explores how a hypothetical humanitarian relief effort can exploit information published on SM sites to gain information relevant to its mission.

This aforementioned situation designates the need for newly developed analytics incorporating credibility analysis to extract reliable information about a conflict event. The next section introduces a SM content credibility analysis approach and prototype fact-finding technology applied to the ongoing conflict situation in Syria.

4. RELIABLE FACT-FINDING FROM SOCIAL MEDIA

Investing in the development of technologies for large-scale event analysis informed by SM content could enable careful observation of events in ongoing and future trouble zones for purposes of ensuring safety or understanding the nature of hostility. However, a key research challenge posed by the Army's prospective use of SM data lies in the need to ascertain data correctness. This problem is often called data *credibility analysis*. There are two general directions in credibility analysis in current research literature:

- The first direction attempts to model how humans evaluate credibility of data. Machine learning approaches are used to train classifiers based on many data samples judged (i.e., labeled) by a human observer as credible or not. The goal is to train the classifier to distinguish credible information as might be judged by a human. These classifiers consider a large number of features pertaining to content, sources, frequency of occurrence, and context, and determine which features are more predictive of human credibility judgment.
- The second approach to credibility analysis uses purely statistical techniques borrowed from data fusion literature. These techniques do not interpret content semantics. Rather, they consider the statistical properties of content dissemination. Modeling data sources as unreliable sensors, the question addressed by statistical analysis techniques is to determine the probability that an observation is true, given the set

of sensors that agree or disagree with it. For this approach to work, one needs to determine the reliability of individual sources as well as their dependencies. Such dependencies arise, for example, when one source may be influenced by another, which may lead it to agree with (e.g., re-post) a piece of data without independently verifying its correctness.

Statistical analysis techniques nicely complement human credibility judgment offering conclusions that are unbiased by content semantics and are based on statistical evidence alone. Note that, the goals of the aforementioned two approaches are somewhat orthogonal. For example, consider a situation where, in the aftermath of a hurricane that disrupted power and means of communication, a rumor spread claiming that the Statue of Liberty was toppled off by the wind and fell. Independently of whether or not a human is likely to believe this claim, the observation could either be true or false. One way to distinguish the two cases is by careful analysis of the statistics involved, such as the number of sources making the claim, their inferred reliability (based on other claims they made), and the relations between these sources (e.g., whether they are independent or not). Given different statistical evidence, the conclusions regarding credibility may differ. Such statistical evidence is hard for a human to mathematically process. Humans are likely to base their judgment on other more semantically-oriented features. It is this complementary nature of statistical analysis that makes it of great value as an aid to analysts, exploiting the superior mathematical processing capabilities of machines in processing large data sets.

4.1 Statistical Credibility Analysis

Recent research on statistical analysis techniques formulated credibility analysis as a maximum likelihood estimation problem, where both the reliability of sources and the credibility of their claims are determined jointly, even if knowing neither in advance. The estimator takes into account source dependencies inferred from patterns of correlations between outputs of different sources. Those source dependencies may be topic-specific and sentiment-specific, as sources will tend to propagate information that is consistent with their own biases and beliefs, while ignoring other information with higher probability.

By observing which sources propagate which claims, it becomes possible to determine the interests of these sources. One can also uncover the latent dissemination (or influence) backbone in a topic-specific and sentiment-specific fashion. To illustrate, we analyzed microblog posts (tweets) collected from Twitter in the aftermath of the Syrian chemical weapons crisis in August 2013. The tweets were crawled for ten days using the keywords “Syria”, “attack”, “dead” or geo-location originating from 120 miles around Syria. Approximately 205,000 tweets matched the query. Table 1 details the crawler information for collection of the tweets.

Table 1.Crawler information for collection of the Syrian chemical attack tweets.

Crawler Information
Longitude: 36.93097
Latitude: 34.78036
Radius: 120
Keywords: syria, attack, dead
Date Range: August 22 – 31, 2013
Size: 805 MB
Number of Cascades: 128K Cascades
Number of Tweets: 205K

During the crisis, disturbing news and images were uploaded to SM in the early morning of August 21st, showing significant numbers of casualties with signs consistent with neurotoxin use, asking for medical supplies and other humanitarian relief. Immediately following these posts, different camps developed different widely circulated

hypotheses regarding what happened, blaming different parties for the deaths. These included a hypothesis that Syrian rebels accidentally detonated chemical weapons while transferring them to another location, a hypothesis that the Syrian government ordered those bombs, and a hypothesis that a third (foreign) party carried out the attack to frame Syria. In our analysis, we lumped all hypotheses into a *neutral* category, a *pro-government*, and an *anti-government* category (referring to the government of Syria). We then manually annotated the 1000 most widely circulated tweets accordingly. Figure 1-a and 1-b show the probability distribution of the sources and claims of these tweets.

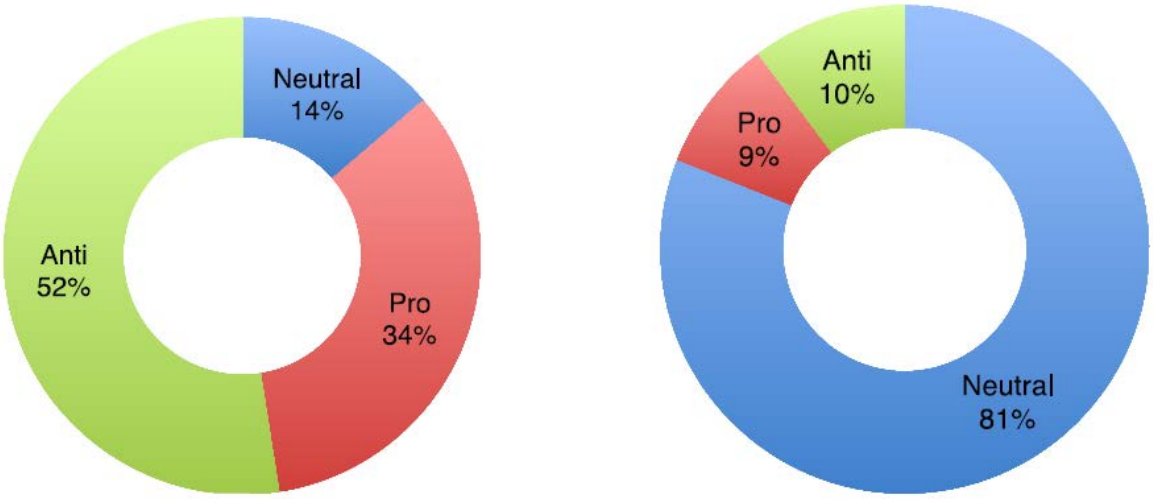

(a) Distribution of source sentiment.

(b) Distribution of claim sentiment.

Figure 1. Source and claim sentiment breakdown of the Syrian August 2013 tweets.

Figure 1 shows that the majority of sources take sides, although they account for only a minority of unique claims. A small number of sources are neutral and make up the majority of unique claims. The fact that a large fraction of sources collectively account for a minority of unique claims arises because polarized claims tend to be largely retweeted. Hence, while there are many participants to the retweet cascades, the number of distinct cascades is small.

Figure 2 shows the probability distribution of the pro-government sentiment. Specifically, we plot the probability that a source posts a pro-government tweet. It can be seen that the distribution is bi-modal, indicating that sources are divided between those who almost always tweet pro-government and those that never do. A small fraction of sources are apparently neutral, as they tend to propagate tweets of both types with roughly the same probability.

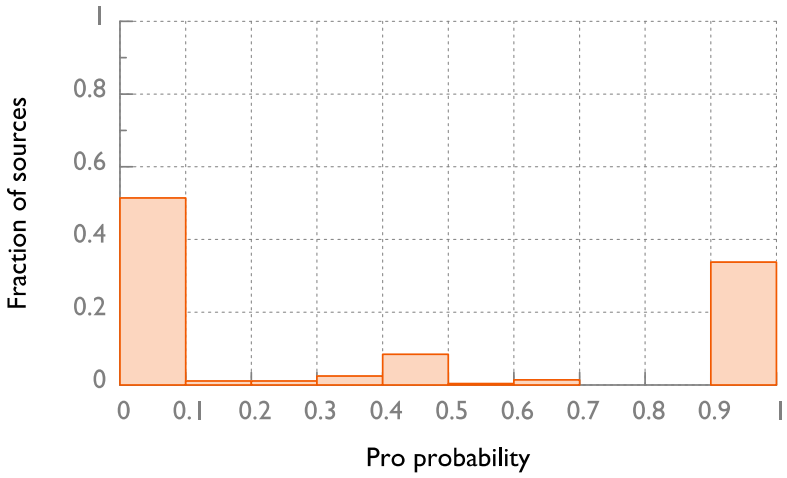

Figure 2. Probability distribution of posting pro-government tweets.

Figure 2 suggests that it is possible to distinguish between a pro-government information dissemination network, as well as, an anti-government network. Those networks map the dependencies between sources when propagating certain types of information. Such dependencies must be considered in statistical analysis. Specifically, they are modeled as correlations, affecting the calculation of probability of having a number of sources make a correlated error (specifically, it reduces this probability). Figure 3 shows the pro-government and anti-government networks. As might be expected, the networks are largely non-overlapping. It is also interesting to note that some nodes seem to follow tweets of the opposite polarity then disseminate their own bias.

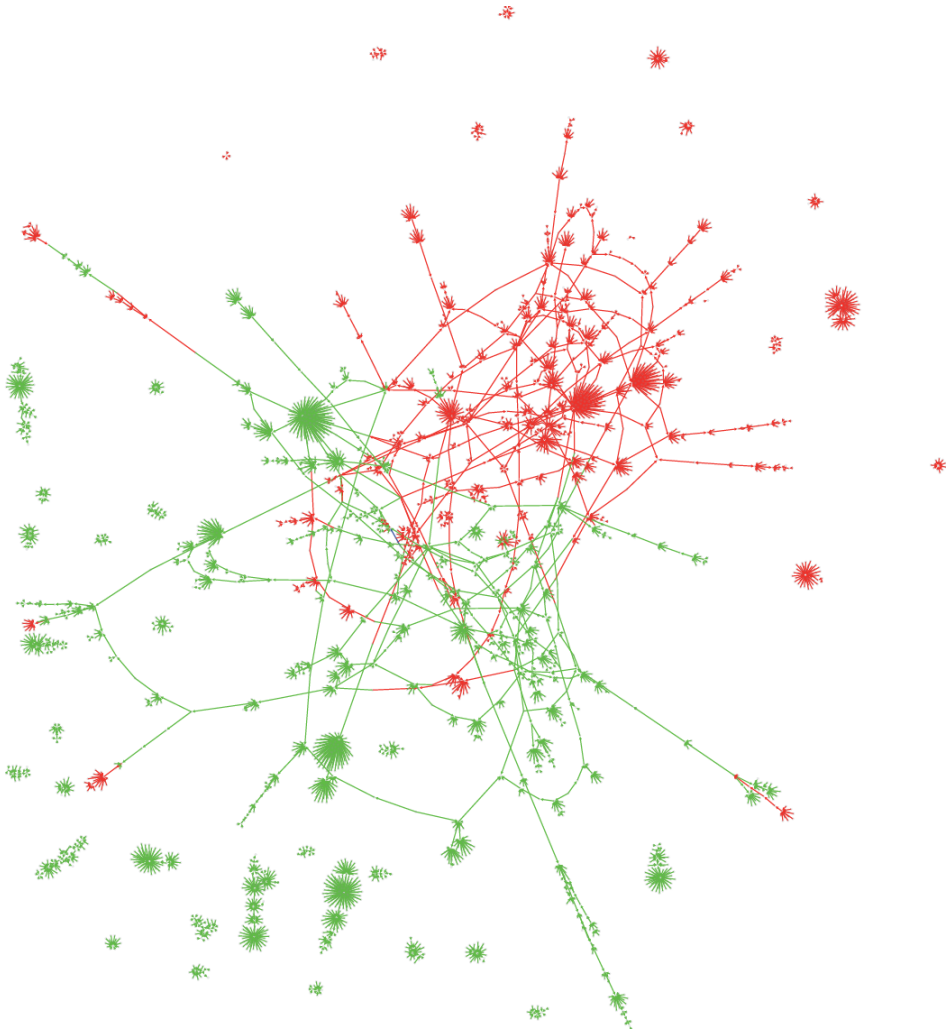

Figure 3. Two extracted latent networks of opposite polarities (green, anti-government; red, pro-government).

For clarity, only a subset of tweet cascades is shown in Figure 3. A tweet cascade is defined as an individual piece of information propagating through the social network. We found that the crawled tweets resulted in 128K cascades. The largest 1000 cascades were selected that consists of 30K tweets and 23K sources. We manually annotated the claims expressed by these cascades among the classes {Pro, Anti, Neutral} based on their polarity being "Pro-government" or "Anti-government". We found that 88 claims were pro-government, 102 were anti-government, and the rest were neutral. We estimated the social network for each of these classes. The pro-government network had around 2K links, the anti-government network had 3K links, and the neutral network had 20K links. Table 2 details the characteristics of the 1000 tweet cascades. These networks, along with the source-claims information (which source made which claim) were used as input to the Apollo Fact-finder.

Table 2. Statistics of the 1000 cascades used in the experiment.

Characteristics of Top 1000 Tweet Cascades
Cascades: 100
Pro Claims: 88
Anti Claims: 102
Neutral Claims: 810
Number of Sources: 22,745
Number of Tweets: 29,216
Pro Tweets: 2,300
Anti Tweets: 3,379
Neutral Tweets: 23,537
Source-claim Edges (Total): 27,498
Source-claim Edges (Pro): 2,031
Source-claim Edges (Anti): 3,123
Source-claim Edges (Neutral): 22,344
Pro Network Edges: 1,934
Anti Network Edges: 2,933
Neutral Network Edges:20,200
Combined Network Edge: 24,738

4.2 The Apollo Fact-finder

Taking the aforementioned information dissemination patterns into account, as well as which sources make which claims, a prototype called the "Apollo Fact-finder" was developed to perform statistical analysis that yields an estimate of the reliability of individual sources and the credibility of individual claims. Statistical analysis in Apollo *does not interpret the text of the tweets*. By default, it also *does not incorporate prior knowledge about the sources*. Such knowledge can be added when available, but it is not a necessary prerequisite for the analysis. The goal is enable analysis even when sources are unknown. The output of the analysis is a triage of the tweets. Given hundreds of thousands or millions of tweets as input, Apollo produces a listing of most statistically important tweets that should be looked at first.

Table 3 shows sample outputs of the statistical analysis. The left column shows examples of tweets deemed credible by Apollo. The right column shows examples deemed less credible or less important based on the statistics of their sources and dissemination patterns. These statistics take into account the latent dissemination network topology, the sources involved, and the inferred reliability of these sources according to the underlying maximum likelihood estimator.

Table 3. Example Apollo Fact-finder results of data triage using statistical analysis.

Triage Result: Recommended for Viewing	**Triage Result: Dismissed/Unimportant**
Medecins Sans Frontieres says it treated about 3,600 patients with 'neurotoxic symptoms' in Syria, of whom 355 died http://t.co/eHWY77jdS0	So sad. All but one of the activists who filmed the chemical attack in Syria died of toxins: http://t.co/7Xc9u8achL
Weapons expert says #Syria footage of alleged chemical attack "difficult to fake" http://t.co/zfDMujaCTV	Saudis offer Russia secret oil deal if it drops Syria via @Telegraph http://t.co/iOutxSiaRs
U.N. experts in Syria to visit site of poison gas attack http://t.co/jol8OlFxnf via @reuters #PJNET	Putin Orders Massive Strike Against Saudi Arabia If West Attacks Syria http://t.co/SFLJ9ghwbt
Syria Gas Attack: 'My Eyes Were On Fire' http://t.co/z76MiHj0Em	Miley Cyrus twerks meanwhile in other news the U.S.A. might declare war on Syria....
Long-term nerve damage feared after Syria chemical attack http://t.co/8vw7BiOxQR	I posted a new photo to Facebook http://t.co/FRWBFC0vKb
Syrian official blames rebels for deadly attack http://t.co/76ncmy4eqb	Two Minds on Syria http://t.co/ogDjKFH7Rs via @NewYorker
Assad regime responsible for Syrian chemical attack, says UK government http://t.co/pMZ5z7CsNZ	We may be going to war in Syria, and somehow Miley Cyrus Is trending on twitter
US forces move closer to Syria as options weighed: WASHINGTON (AP) — U.S. naval forces are moving closer to Sy... http://t.co/F6UAAXLa2M	Syrian Chemical Weapons Attack Carried Out by Rebels, Says UN (UPDATE) http://t.co/lN4CkUePUj #Syria http://t.co/tTorVFUfZF
400 tonnes of arms sent into #Syria through Turkey to boost Syria rebels after CW attack in Damascus --> http://t.co/KLwESYChCc	For those in the US, please text SYRIA to 864233 to donate $10 via @unicefusa http://t.co/YMXnrk1jcb #childrenofsyria
UN Syria team departs hotel as Assad denies attack http://t.co/O3SqPoiq0x	Attack! http://t.co/wY5KKm7R3s
Vehicle of @UN #Syria #ChemicalWeapons team hit by sniper fire. Team replacing vehicle & then returning to area.	A fathers last words to his dead daughters killed by Bashar al-Assad & his supporter army with chemical weapon attack http://t.co/DN25pLfCq8
International weapons experts leave Syria, U.S. prepares attack. More @ http://t.co/4Z62RhQKOE	What the media isn't telling you about the Syrian chemical attack http://t.co/LQ479S1Tiv
Military strike on Syria would cause retaliatory attack on Israel, Iran declares http://t.co/M950o5VcgW	France on the phone. Apparently they surrendered to #Syria weeks ago.
Asia markets fall on Syria concerns: Asian stocks fall, extending a global market sell-off sparked by growing ... http://t.co/06A9h2xCnJ	Poll: Do you think the chemical attack in #Syria could have been a false flag attack to push for war? RT for yes. Favourite for no
UK Prime Minister Cameron loses Syria war vote (from @AP) http://t.co/UlFF1wY9gx	Lebanon was once part of Syria and will forever be with Syria. #PrayForSyria #PrayForLebanon

The analytic foundations and principle of operation behind Apollo were described at length in prior publications[14,15]. This research prototype presents a query API to the user that allows specifying regions of interest and supplying filtering criteria (keywords) for tweets originating from those regions. Each query creates a task that collects, in real time, a stream of tweets matching the user-specified location and keyword preferences. This real time stream is processed and ranked by the triage engine, yielding an output that allows the user to navigate the evolving event timeline, and view the key highlights of an event at a specified time. An example query is shown in Figure 4 (top). The resulting timeline is shown in Figure 4 (bottom). The user is allowed to scroll back and forth in time to view the key recommended tweets from different time frames.

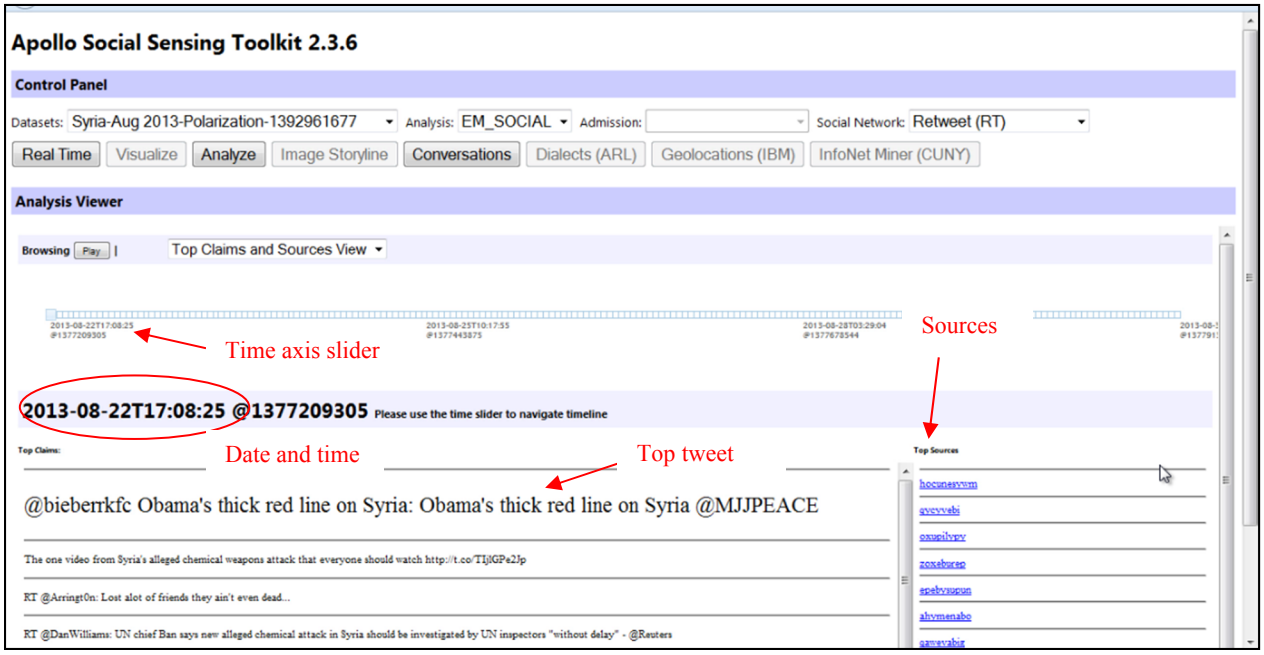

Figure 4. Apollo Fact-finder API: (a) Top, an example query. (b) Bottom, the output timeline of the analysis algorithm.

5. DISCUSSION

An important open challenge is to combine the statistical analysis techniques presented above with machine learning approaches that mimic human credibility judgment. Detailed inspection of our results suggests that tweets with insufficient support were not picked up by statistical analysis, as one might expect. These tweets may not necessarily be wrong. In fact, they jointly constitute a large majority of input. Exploiting the features learned by classifiers that mimic human judgment in the analysis of such tweets could allow discovery of important but low-profile events. Conversely,

human judgment may occasionally err in situations where unexpected events occur. Observing disagreement with significant statistical evidence can help catch such judgment errors.

For statistical analysis techniques to work, the analysis algorithm has to understand which tweets make the same statement, such that they are clustered together and viewed as the same claim. This clustering allows support for different claims to be properly computed. Clearly, there may be several different ways to express a claim. In the current version of Apollo, a "distance function" is used that accepts pairs of tweets as input and yields a degree of similarity as output. The distance function allows the body of incoming tweets to be clustered by similarity.

Motivated by the need to process tweets in arbitrary languages, no natural language processing is performed inside Apollo's distance function. Rather, tweets are viewed as sets of abstract tokens (words). Tweets with a large degree of overlap between their sets of tokens are grouped together as the same claim. Research is needed on the right degree of semantic inspection that is required to make a proper similarity judgment. For example, inserting a "not" somewhere in a long sentence can reverse its meaning. Hence, structurally similar tweets can denote different things. Similarly, use of different synonyms can lead Apollo to incorrectly separate semantically similar tweets into different claims. While Table 3 suggests that the current function already leads to a meaningful tweet categorization, a minimum level of semantic inspection can be incorporated to catch the majority of such cases and hence improve outcome.

Algorithm scalability must keep pace with the exponential expansion of SM use for information sharing and influence activities. The emerging capabilities provided by Apollo and similar exploitation tools provide a triage method to highlight information with potential highest-value for event detection. Additional automated capabilities are in development and must be delivered quickly for Army analysts to realize utility from SM products. The value of SM, apart from traditional physical-sensing based sources, lies in the ability to sense new insights into how attitudes are changing over time in a population where Army sources may be few and the analyst unfamiliar with cultural norms and behaviors. The structure of SM inputs allows trending of sentiment such as the use of particular terminology or hashtags that could indicate attitudinal change towards more sectarian or more tolerant opinions, with key inflection points identified and matched with possible real-world drivers. The rapid triaging of popular characteristics in a population segment, such as discussion topics, sentiment toward leaders and major political ideas, sensitive locations and social groups, and calls for social action, all are important drivers that demand Army notice and monitoring in locations that are unstable and might demand military action in the future.

Finally, in this paper, when exploring polarization, tweets were manually annotated into the different categories of interest (such as pro-government and anti-government). An interesting research topic is to exploit the natural clustering of source and tweets in an automated fashion. Since we observed that sources of different biases have largely different forwarding patterns for different categories of tweets, an interesting question is whether the categories of tweets and sources can be uncovered in an automated fashion, solely based on source forwarding behavior. Such automation would then remove the need for manual annotation.

Polarization analysis can be of great use to decision makers who face daily decisions on when and how to execute operations of various types. Social action by crowds as witnessed in Syria and other nations since the advent of SM is extremely dynamic and difficult to predict. This makes course of action planning and execution even more difficult because action can make a bad situation worse. In the case of Syria, polarization analysis has the potential to identify and evaluate online statements and trends to real-world developments. For example, did the attention paid in SM to Syrian rebel groups translate into material support for those groups? And, if so, was that support a direct result of the SM attention? Do more polarized online communities necessarily mean more divisions on the ground? Given that polarized political views are destabilizing to open and fair democracies, how can national and international leaders use SM to influence more moderate views by local citizens or improve humanitarian conditions in the conflict area? What political, economic, and social actions can be taken by the national and international communities to reduce polarization and support stable civil life?

CONCLUSION

Syria's socially mediated violent conflict is likely to be a model for future crises. The paper suggests ways in which SM could be used to better understand what happened. During a conflict it is especially important to understand how information is produced, how it flows through social networks, and how it gains or loses credibility and significance

with relevant external audiences. New inventive methods of credibility analysis are needed to benefit Army operations by exploiting the wealth of information about a conflict event circulating online while addressing the serious challenges posed by inaccurate, unrepresentative, and deliberately bias SM content.

The paper explores polarization results from a statistical credibility analysis and introduces the prototype Apollo Fact-finding technology that formulates both the reliability of sources and credibility of claims from "big data" streams. These analysis techniques were applied to tweets collected in the aftermath of the Syrian chemical weapons crisis in August 2013. Developing methods for authenticating SM information is extremely important from a command and tactical level of combat in situations where conflict conditions limit access to the battlefield and deployment of assets is not possible.

ACKNOWLEDGEMENTS

Research reported in this paper was sponsored by the Army Research Laboratory and DTRA grant HDTRA1-10-1-0120, and was accomplished under Cooperative Agreement Number W911NF09-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

REFERENCES

[1] Tsukayama, H., "Twitter turns 7: Users send over 400 million tweets per day," (21 March 2013) http://articles.washingtonpost.com/2013-03-21/business/37889387_1_tweets-jack-dorsey-twitter

[2] Eldon, E., "Facebook climbs toward 700 million users worldwide, with steady growth in the US," (21 April 2011) http://www.insidefacebook.com/2011/04/21/facebook-climbs-toward-700-million-users-worldwide-with-steady-growth-in-the-us/

[3] Rogers, E. M., [Diffusion of Innovation], Free Press, 5th ed. (2003).

[4] Kwak, H., Lee, C., Park, H., and Moon, S., "What is Twitter, a social network or a news media?" WWW 2010, April 26-30, (2010).

[5] Sohail, R. M. and Chebib, N., "The reasons social media contributed to 2011 Egyptian revolution," International Journal of Business and Research Management, 139-162 (2011).

[6] Serena, C., Tingstad, A., Winkelman, Z., Tkacheva, O., Bim, M., Marcellion, W., Mobley, B., Clarke, C. P. and Paul, C., "The role of social media in future army operations," RAND Arroyo Center Force Development and Technology Program, Draft Report (2013).

[7] Mayfield, T. D., "A commander's strategy for social media," Joint Forces Quarterly 60(1), 79 (2011).

[8] Associated Press. "Tweets, Pics Give Real-time Peek into North Korea." (12 July 2013) http://bigstory.ap.org/article/tweet-allow-peek-life-nkorea-real-time

[9] Zeng, D., Chen, H, Lusch, R. and Li, S., "Social media analytics and intelligence," Intelligent Systems, IEEE 25(6), 13-16 (2010).

[10] Goolsby, R., "Social media as crisis platform: The future of community maps/crisis maps," ACM Transactions on Intelligent Systems and Technology (TIST) 1(1), 7 (2010).

[11] Comninos, A. "Twitter revolutions and cyber crackdowns," Academia.edu. (12, June 2013) http://academia.edu/633706/Twitter_revolutions_and_cyber_crackdowns_User-generated_content_and_social-networking_in_the_Arab_spring_and_beyond

[12] Mitchell, A. and Hitlin, P., "Twitter Reaction to Events Often at Odds with Overall Public Opinion," Pew Research Center, (2013).

[13] Lynch, M., Freelon, D. and Aday, S., "Syria's socially mediated civil war." United States Institute of Peace, 91, (2014).

[14] Wang, D., Amin, T., Li, S., Abdelzaher, T., Kaplan, L., Gu, S., Pan, C., Liu, H., Aggrawal, C., Ganti, R., Wang, X., Mohapatra, P., Szymanski, B. and Le, H., "Humans as sensors: An estimation theoretic perspective," 13th International Conference on Information Processing in Sensor Networks (ACM/IEEE IPSN), Berlin, Germany, (2014).

[15] Wang, D., Le, H., Kaplan, L. and Abdelzaher, T., "On truth discovery in social sensing: A maximum likelihood estimation approach," 11th ACM/IEEE Conference on Information Processing in Sensor Networks (IPSN), (2012).

1 (PDF)	DEFENSE TECH INFO CTR ATTN DTIC OCA (PDF)
2 (PDF)	US ARMY RSRCH LABORATORY ATTN IMAL HRA MAIL & RECORDS MGMT ATTN RDRL CIO LL TECHL LIB
1 (PDF)	GOVT PRNTG OFC ATTN A MALHOTRA
6 (PDF)	U.S. ARMY RSRCH LAB ATTN RDRL CII C E K BOWMAN ATTN RDRL CII C M THOMAS ATTN RDRL CII C S KASE ATTN RDRL CII C T HANRATTY ATTN RDRL CII B BROOME ATTN RDRL CII T V M HOLLAND

INTENTIONALLY LEFT BLANK.